Operating System

Active Directory Interoperability and Metadirectory Overview

Strategy White Paper

Abstract

Identity is the summary of information about people, applications, and resources scattered in directories and databases throughout most IT enterprises. This paper addresses solution requirements, using Microsoft® Windows® 2000 and the Active Directory™ directory service, for dealing with disparate identity information, including the sharing of identity information between different resources, the distribution of identity changes amongst various resources, and ensuring that related data remain consistent throughout the enterprise.

© 1999 Microsoft Corporation. All rights reserved.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

Microsoft, Active Directory, the BackOffice logo, Visual Basic, , Windows, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Other product or company names mentioned herein may be the trademarks of their respective owners.

Microsoft Corporation • One Microsoft Way • Redmond, WA 98052-6399 • USA

0798

CONTENTS

OVERVIEW....1
Customer Situation 1
Solution Requirements 1
Solution Alternatives 1
Enterprise Identity Management with Active Directory 2

ENTERPRISE IDENTITY MANAGEMENT....3
The Identity Management Challenge 3

SOLUTION REQUIREMENTS....4
Connectivity 4
Brokering 4
Change Event Processing 5
Data Aggregation Capabilities 5
Related Object Tracking 6
Integrity Management 6
Ownership 7
Failure Management 7
Referential Integrity 8

SOLUTION ALTERNATIVES....9
Multi-Directory Access 9
Synchronization Connectors 10
LDAP Proxy Interfaces 10
Hub and Spoke Architectures 11
Directory Consolidation 12

MICROSOFT SOLUTIONS....13
Multi-Directory Access: Active Directory Service Interfaces 13
Synchronization Connectors: ADC and DirSynch 13
LDAP Proxy Interfaces 14
Hub and Spoke Architectures 14
Directory Consolidation 15

Conclusion....17

OVERVIEW

Customer Situation

Information about people, applications, and resources is scattered throughout most IT enterprises, and is continuing to proliferate. An increasing amount of this *identity data* is stored in standards-based directory services, but the majority remains stored in databases and other specialized forms. Customers confront identity management challenges in many different forms:

- **Single sign-on initiatives.** Manage username, password, and access rights information across many different platforms and applications.
- **Global address book applications.** Synchronize mailbox information between different e-mail directories used within a company.
- **“Hire & fire” solutions.** Quickly propagate information about newly hired employees to all systems that require identity data, and quickly perform the same processes in reverse when employees leave.
- **E-commerce applications.** Synchronize information such as digital certificates for suppliers and extranet users with e-commerce directories that reside outside of firewalls.

With each additional application and platform that a customer deploys, the number of places where they must manage identity data increases. This forces companies to manage a significant amount of duplicated and related information in many different places.

Solution Requirements

Theoretically, the simplest solution is to have a single enterprise directory that holds all information about users, machines, networks, and applications in the company. For many reasons, including political boundaries, this goal will not be achieved quickly–if ever—at most companies. It is this reality that leads companies to look for solutions that link different directory services and applications together and provide a consistent way to store, access, and manage their identity data. Because this assumes that identity data will continue to exist in many places, solutions must provide:

- **Connectivity,** to enable the sharing of identity information between many different directory services, databases, and applications.
- **Brokering functionality**, to distribute changes made in one directory or application to other identity repositories in the enterprise affected by the change.
- **Integrity mechanisms,** to ensure that related data remains consistent throughout the enterprise and observes ownership and referential integrity rules.

Solution Alternatives

It is unlikely that there will be one solution to all identity management challenges because the issues are so varied. Rather, the average company should look to deploying several different approaches, including:

- **Multi-directory access technologies,** to simplify writing applications and administrative scripts by providing developers with a single programming interface to multiple directory service and database technologies.
- **Synchronization connectors,** to simplify management by keeping pairs of directory services synchronized with each other automatically.
- **Meta-directories technologies,** to provide companies with advanced connectivity, brokering, and data integrity management capabilities.
- **Directory consolidation strategies,** to enable companies, over time, to reduce the total number of directories that they need to manage.

Enterprise Identity Management with Active Directory

Microsoft recognizes the importance of solving identity management challenges and is delivering a comprehensive set of solutions based on the Active Directory directory service of the Windows 2000 Server operating system.

- **Active Directory Service Interfaces (ADSI).** ADSI is a set of extensible, easy-to-use directory programming interfaces based on Microsoft's Component Object Model (COM).
- **Active Directory Connector (ADC) and Microsoft Directory Synchronization Services (MSDSS).** ADC provides drop-in synchronization services between Active Directory and Microsoft Exchange Server versions 5.0 and 5.5, enabling administrators to eliminate redundant administrative steps. ADC also provides support for other e-mail and LDAP-compliant directories. MSDSS is a component of Microsoft's Services fo Netware v.5 product that makes it easy for administrators to synchronize changes between Active Directory and Novell's NDS.
- **Meta-directory technologies.** By integrating and enhancing meta-directory technologies gained from Microsoft's recent acquisition of Zoomit Corporation, Microsoft will deliver a comprehensive set of meta-directory and enterprise identity management technologies based on Active Directory.
- **Directory consolidation opportunities.** Microsoft is updating many of its products, such as Exchange Server, to use Active

Directory instead of separate repositories and is working with leading independent software vendors to do the same.

Based on the breadth and depth of these solutions, Microsoft believes that Active Directory is the ideal platform on which to implement enterprise identity management solutions.

Identity is the summary of information about people, applications, and resources scattered in directories and databases throughout most IT enterprises. Examples of identity data associated with people include name, mailbox, salary, and job title information. Application identity information includes network addresses where

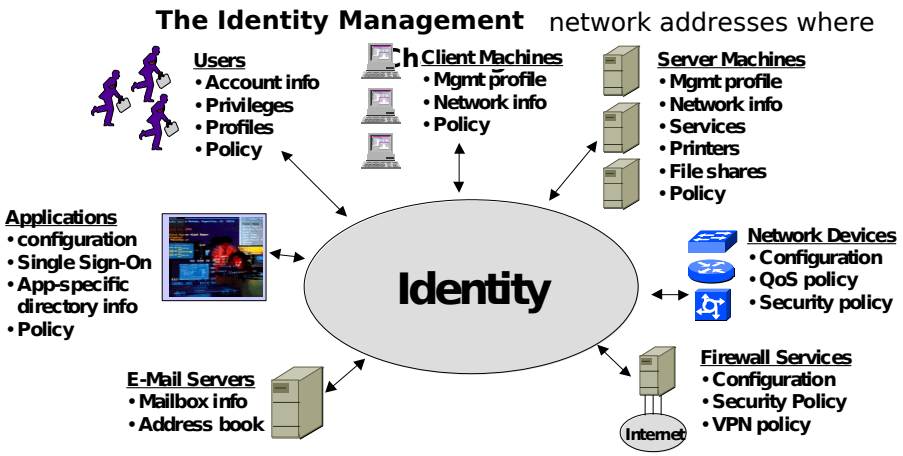

clients can find servers and lists of services that applications can provide. Network resources, such as printers, also have identity attributes; for instance, their location and the printing capabilities they support.

The Identity Management Challenge

The diversity of identity data and number of places where identity resides raise a number of management challenges:

- Not all identity data is kept in directories or exposed through a directory service interface such as LDAP. For example, many systems only expose identity information through specialized application programming interfaces (APIs).
- Identity information frequently is duplicated in multiple places and tends to drift out of synchronization over time if left unchecked.
- Typically, there is no single place where administrators and applications can go to access or manage an aggregated view (sometimes called a ‘join’) of an enterprise’s identity information.
- The number of places where companies must manage identity data increases with each additional application and platform they deploy.

SOLUTION REQUIREMENTS

These challenges make it difficult for companies to implement comprehensive and integrated identity management solutions, and translate directly to cost and complexity.

In the past, many companies have tried to create a single directory to hold all enterprise identity information. Most of these efforts failed for several simple reasons:

- Many applications cannot be modified easily to use directories.

Connectivity Requirements

- There are good reasons, such as advanced security requirements, why some applications need to keep identity in their own formats.
- Political boundaries inhibit complete consolidation regardless of what is technically possible.

This suggests that identity data will continue to exist in many places and that companies need to find ways to make different directory services and application repositories work together. Assuming that there will be many identity repositories, solutions must provide connectivity to many forms of identity data, brokering functionality to manage the flow of

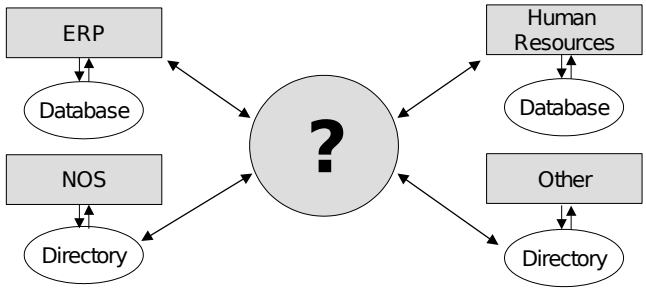

information between repositories and mechanisms for maintaining data integrity throughout the identity management infrastructure.

Connectivity

Connectivity requirements are simple: the more directory services, databases, and applications to which identity management solutions can connect, the more value they can offer. To say that an identity management solution can connect to a given repository, it must be able to:

- Obtain information about what has changed in the repository
- Add new objects to the repository
- Delete objects from the repository
- Change an existing object’s attributes to different values

To be a comprehensive solution, technologies should be able to connect to data in:

- Standards-based directory services via LDAP version 3
- Databases via access methods such as SQL
- Applications where the only interface to identity information is through application programming interfaces (API) and no directory interface is available.

Change Event Processing

Brokering

Brokering is the process of managing the flow of identity information between repositories. Brokering functionality must be able to:

- Detect changes to identity data and propagate updates to other repositories

Data Aggregation into a Meta-Directory

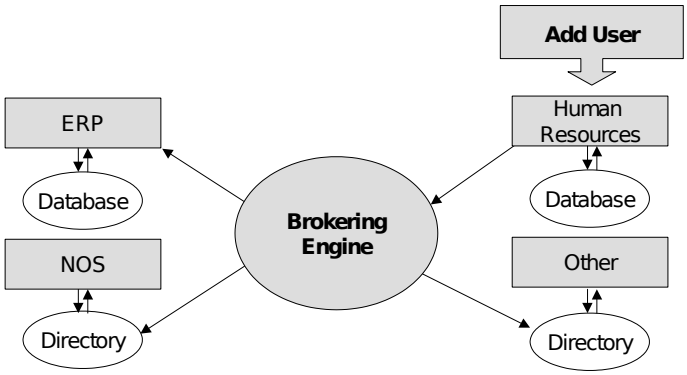

- Aggregate data from different repositories into meta-directories that contain a holistic view of identity data from across the enterprise
- Track related objects as they change their positions in directory trees and other repositories due to periodic reorganization.

Change Event Processing

Change events occur any time administrators, users, or applications add, delete, or modify a piece of identity data in a repository. Without the ability to detect and process changes, identity data quickly becomes disorganized. Identity management solutions therefore must provide features to detect changes, perform necessary data format translations,

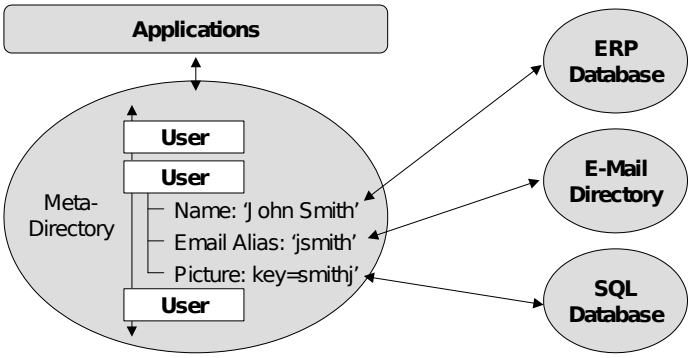

and then cause related updates to occur in all repositories that should reflect the change. For example, if an administrator adds a person to the Human Resources database because they have joined the company, this change event needs to cause systems that the person will use to reflect the addition.

Data Aggregation Capabilities

While identity information resides throughout most enterprises, *meta-*

Tracking Related Objects

directories that contain an aggregation of identity data from many other repositories can offer great value. For example, applications can access a variety of information in one place using a single access method and security model instead of having to interact will all of the source repositories directly. Meta-directories also maximize performance because data can be stored in indexed form and there is no need to fetch data from sources, which may reside across wide area network (WAN) connections, at run time. To offer the greatest value, data aggregation capabilities must be able to:

- Gather and incorporate information from many sources including directories, databases, and applications.
- Group related information together even though it may be stored in different ways in different places. For example, a user named John Smith might have data stored under names such as *John Smith*, *jsmith* and *smithj* in different systems.
- Push changes back out to sources when users or applications make

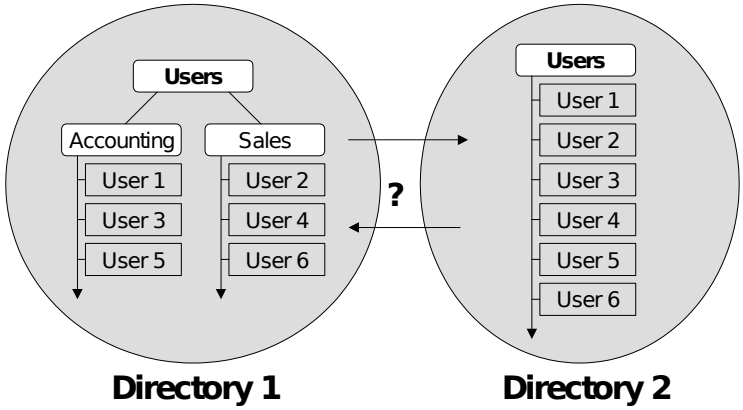

changes to the aggregated view. This means that meta-directories must be integrated with change event processing infrastructures.

Related Object Tracking

When administrators initially deploy identity management solutions, they must be able to provide ways for brokering features to establish relationships between related pieces of identity data stored in different

repositories. For example, administrators must be able to tell the broker that *John Smith*, *jsmith* and *smithj* are all the same person. Then, the broker must be able to track relationships as identity data is reorganized periodically. For example, solutions must not lose track of users simply because they change positions in a directory tree structure due to a move from the Accounting department to the Sales group.

Managing Ownership Relationships

Integrity Management

Integrity management is the process of ensuring that identity data does not become corrupt or out of synchronization between repositories as changes and brokering processes occur. Integrity management functionality must be able to:

- Maintain identity data ownership relationships
- Act appropriately when failures occur
- Maintain referential integrity between identity data

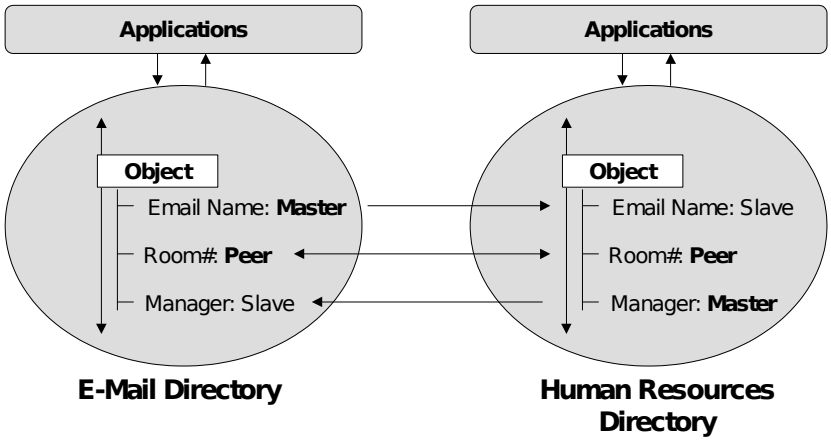

Ownership

An important aspect of enterprise identity management is recognizing that there are important ownership relationships that must be maintained between applications and data. For example, a person's mailbox name is owned by the e-mail system that hosts the mailbox. Within most companies, the human resources (HR) system owns the data corresponding to whether or not a person is an active employee. With no enterprise identity management infrastructure in place, these ownership relationships are preserved by default because no other applications have the ability to access and update e-mail and HR data. With synchronization connectors and brokers deployed, however, the situation changes.

Consider the case where mailbox information is being synchronized to the human resources (HR) directory by a connector. If the connector is not configured correctly, a user could change the mailbox attribute in the HR system and the connector would overwrite the mailbox value in the e-mail directory – causing tremendous confusion.

Managing Failures and Maintaining Referential Integrity

Solving the problem is not as simple as just preventing changes from flowing backwards to the e-mail directory. The HR system may be the owner of information—such as the name of a person's manager—that needs to be flowed back to the directory in the e-mail system. Other attributes, such as a person's office number, may have no clearly defined ownership and should be data that anyone can update.

As a solution requirement, connectors and brokers must enable administrators to define and enforce ownership relationships at the attribute level. In this way, if a change observes ownership rules, connectors or brokers allow it to pass through. If an administrator mistakenly changes a mailbox attribute in the HR directory, the identity management solution would simply set the attribute back to the value contained in the e-mail directory.

Failure Management

The ability to propagate a change to multiple repositories is a key requirement for brokering technologies. Yet, any time a broker makes multiple updates, the opportunity exists for one or more of the updates to

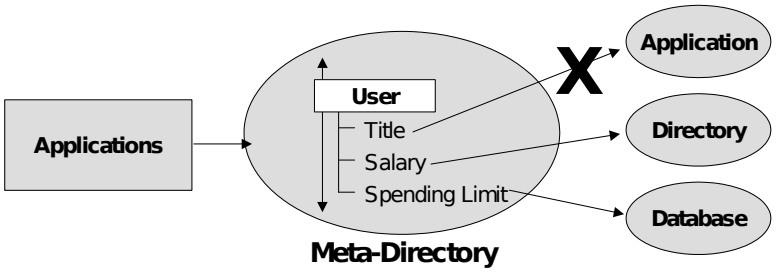

fail and data in different repositories to become inconsistent. For example, if a person is added to the HR system and the broker is unable to contact the e-mail system to add a mailbox, identity data will be left in a state of confusion. Typically, this means that an administrator will have to investigate the situation and make corrections by hand.

In database systems, this challenge is usually addressed with mechanisms such as transactions that ensure that all updates occur successfully or are rolled back as a unit. Unfortunately, most directory services and application programming interfaces do not support transactions. This means that identity management solutions must find other ways, such as using log-based *desired state* mechanisms that continue to request changes until confirmed, to ensure that all repositories eventually reflect changes requested by brokers.

Referential Integrity

Another challenge that identity management solutions share with databases is maintaining *referential integrity* between repositories. Referential integrity refers to the need to maintain relationships between the values of related pieces of data in different locations. For example,

identity management solutions must to be able to ensure that a person's title listed in the human resources system is consistent with their spending limit in the procurement system. Databases help solve this

Multi-Directory Access
An Elusive Goal

challenge by providing stored procedure and trigger features that enable administrators to execute a business rule each time a data value

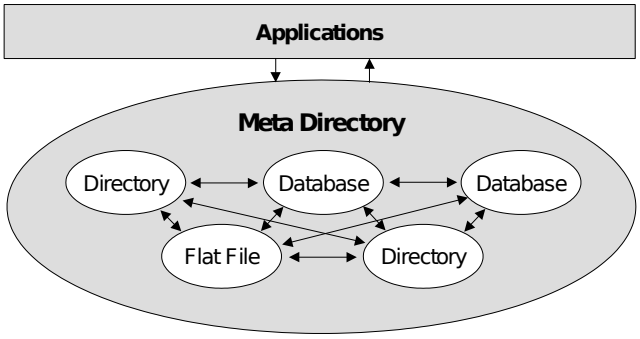

changes. As directory services do not provide similar features today, identity management solutions must provide the capability to execute business rules within brokering engines and reject changes that do not meet referential integrity requirements.

It would be convenient if there were a product that could simply encapsulate directories, databases, and popular application repositories

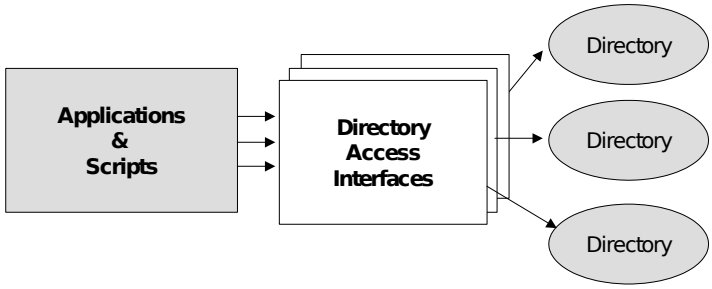

into a 'directory of directories' and handle all of the difficult issues mentioned earlier. Given the complexity of the challenge, however, such a product may never exist. Instead, companies should plan to implement a number of identity management techniques.

Multi-Directory Access

One way to manage multiple sources of identity data is to write applications that interact directly with each of the repositories that hold identity data relevant to a particular management operation. For example, administrators could implement a hire & fire solution by writing a script that, when an employee joins the company, adds appropriate identity information to the HR system, network operating system directories, and each application that the employee will need to use.

The primary advantage of this approach is simplicity. Developers need to learn only one data access paradigm. Applications and scripts can

Synchronization Connectors

contain business rules for updating data and maintaining consistency across repositories. Error checking is straightforward and no synchronization of changes is needed as scripts and applications update all repositories directly.

Tradeoffs of the multi-directory access approach include:

LDAP Proxy Interfaces

- This approach does not create a meta-directory where applications, users, and administrators can go for a holistic view of identity data
- When other applications make changes in individual repositories, there is no mechanism to propagate the changes to other directories,

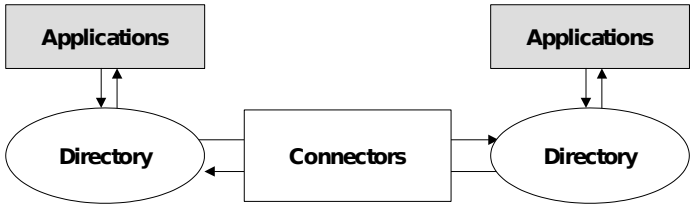

databases, and applications affected by the change.
- There are no automatic mechanisms in this approach for maintaining consistency of related information stored in different repositories.

Synchronization Connectors

Synchronization connectors are applications that know how to recognize changes in one repository and propagate them to another repository. Connectors typically maintain a one-to-one relationship and companies need to have a specific connector for each pair of repositories that they want to manage. Connectors also can be quite versatile, handling situations where one side is a directory and the other is a database or an application.

Synchronization connectors simplify management by enabling administrators to designate certain directories as management focal points where they make changes and rely on connectors to push the changes out to other repositories. Because they require just installation and basic configuration, synchronization connectors also are simple to deploy.

As with multi-directory access techniques, however, connectors do not create meta-directories. Connectors also do not provide a way to implement business rules that may be needed to handle issues such as data transformation and referential integrity of data between

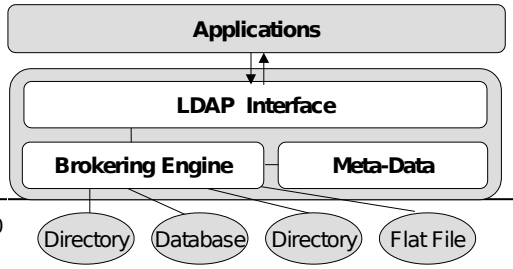

repositories; they excel where relationships between repositories are simple.

LDAP Proxy Interfaces

An *LDAP proxy* is a service that provides an LDAP-compliant interface and knows how to interact with multiple repositories to resolve queries and updates at run time. For example, an LDAP proxy, when presented with a

query request for a user object corresponding to *John Smith*, could retrieve some information from the HR database, some from an operating system directory, and the remainder from an enterprise resource planning (ERP) application. The proxy then returns the data as a single object that appears to have been retrieved from a single source; updates to the data follow a similar path in reverse. Using meta-data repositories, proxies can support a variety of object types.

LDAP proxies have the advantage of being able to provide a single, virtual interface to a wide range of information. By providing update capabilities, proxies also can act as virtual meta-directories. While compelling, proxies also have a number of limitations:

- Because object data is retrieved at run time, performance can suffer. Performance issues get worse when data must be fetched from across WAN connections, when applications query by attributes that are not indexed in one or more of the repositories accessed by the query, and when queries must return large amounts of data. Also, the speed of queries and updates are limited by the slowest resource accessed by the request.
- Proxies provide no ability to synchronize changes made by applications that bypass the proxy interface and access identity repositories directly.
- Because of the wide range of security models used by identity repositories, proxy interfaces typically have to run with administrator

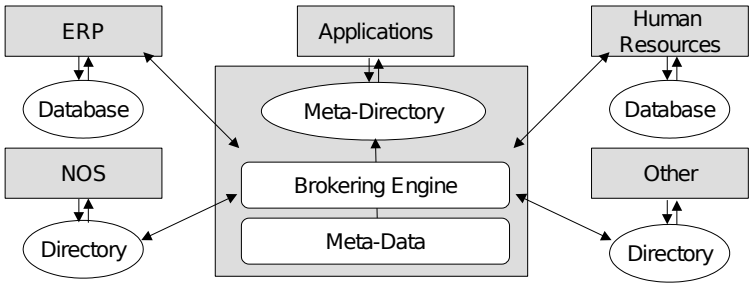

privileges and implement their own access control paradigms.

Hub and Spoke Architectures

Hub and spoke architectures focus on connecting to many repositories,

capturing changes to identity data, and using a brokering and rules engine to propagate the changes to any other repository that should

Directory Consolidation

reflect the change. In this type of architecture, an addition of an employee to the HR system could trigger the broker to create user accounts, set up an e-mail mailbox, and create rights in the ERP system.

There are several advantages to hub and spoke approaches:

- It is easy to create a meta-directory as part of the brokering process and to treat the meta-directory repository as another source of change information. This enables the meta-directory to act as a focal point for queries and updates.
- The brokering engine can handle different data formats, track object location, and maintain attribute ownership relationships. This enables hub and spoke approaches to integrate well with existing business processes.
- Rule processing capabilities enable companies to implement referential integrity protections and add value to the data stored in the meta-directory.

The most significant downside to hub and spoke architectures is that, to obtain maximum business value, companies must develop the rules used

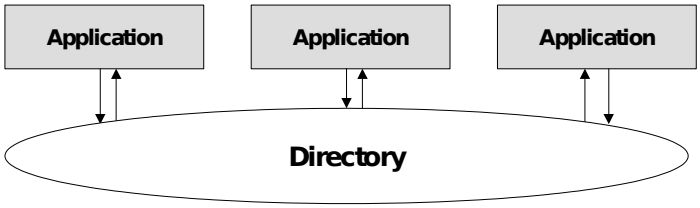

to control data flow in the hub. This downside is offset somewhat by products that offer some form of templates and pre-built rule sets.

Directory Consolidation

The ultimate identity management technique is to implement consolidation directories where multiple applications use the same directory service. While most companies will not get to a single directory service any time soon, each time they can eliminate a directory, the overall level of identity management complexity will fall. For example, within the consolidation directory no synchronization is required and there is a single set of data semantics. When supported by applications, a consolidation approach is also the easiest to implement because there is no need to set up connectors or business rules. The only downside is that, for this approach to work, application developers must provide administrators with the option to use consolidation directories instead of an application-specific repository.

MICROSOFT SOLUTIONS

Microsoft understands the breadth and depth of enterprise identity management challenges and is delivering a comprehensive set of

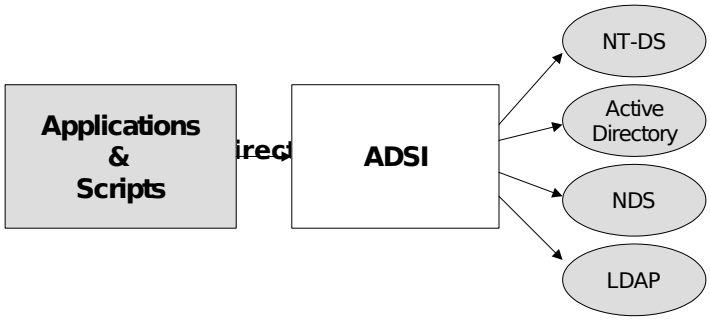

solutions based on the Active Directory service of Windows 2000 Server.

Multi-Directory Access: Active Directory Service Interfaces

To make it easier to write applications and scripts that access a variety of directory services and other identity repositories, Microsoft developed Active Directory Service Interfaces (ADSI). Based on Microsoft's Component Object Model (COM), ADSI provides developers with a set of extensible, easy-to-use programming interfaces to write applications that access and manage:

- The Security Account Manager (SAM) in Windows NT® Server 4.0
- The Active Directory service in Windows 2000 Server
- Novell's NDS
- Any LDAP-based directory

ADSI also provides support classes that simplify common directory programming tasks such as adding new users, managing printers, and locating resources throughout the distributed computing environment. For example, an administrator could write a simple application using the Microsoft Visual Basic® development system (or any tool that supports COM) to reset passwords for all users in a selected container within Active Directory.

To further simplify identity management, Microsoft integrated ADSI with the Microsoft OLE DB common data access framework. This enables millions of developers who understand database-centric client/server programming techniques to write directory-enabled applications with no new learning required. Integration with OLE DB also enables developers to access directory-based data using SQL instead of LDAP and supports advanced features such as heterogeneous 'joins' with SQL data.

Synchronization Connectors: ADC and DirSynch

Microsoft expects to deliver two sets of connector technologies for Active Directory. First, there is the Active Directory Connector (ADC). ADC

ADC and MSDSS synchronizes Active Directory with popular e-mail systems including Exchange Server versions 5.0 and 5.5, Lotus Notes, and Novell's GroupWise. ADC also supports directory services that provide an interface compliant with LDAP version 3. Microsoft expects to ship ADC

with Windows 2000 Server.

Second, Microsoft delivers technologies for managing Active Directory and NDS together in the Services for Netware v.5 (SFNW5) product. In particular, the Microsoft Directory Synchronization Services (MSDSS) component of SFNW5 makes it easy for administrators to synchronize changes between Active Directory and Novell's NDS.These technologies are based on a flexible LDAP-based control called DirSync that enables efficient synchronization of information between heterogeneous directories and Active Directory.

Microsoft also designed the DirSync control to make it easier for developers to build synchronization products that capture changes occurring within Active Directory and propagate them to other directories automatically. Microsoft's design of the DirSync control represents an advance in synchronization technologies:

- DirSynch supports capturing changes at the attribute level, enabling developers to build high-performance connections with Active Directory.
- DirSynch is compatible with the design of most replicated directory services.
- DirSynch allows efficient resynchronization after server failures.

There is already considerable support for DirSync in the industry. For example, a number of leading providers of meta-directory and synchronization products have voiced their intention to use the DirSync control to integrate their products with Active Directory. This ensures that customers will be able to obtain a wide range of synchronization connectors for Active Directory.

LDAP Proxy Interfaces

LDAP proxy interfaces appeal to companies because they promise access to a wide range of identity information in the form of a *virtual replica,* without the overhead of actually replicating data and propagating changes across the network. However, as mentioned earlier, Microsoft

believes that the downsides to LDAP proxy approaches are significant. Therefore, Microsoft will not focus on providing LDAP proxy solutions directly but will work with ISVs that specialize in this area to ensure connectivity to Active Directory.

Hub and Spoke with Zoomit Technologies

Hub and Spoke Architectures

Directory Consolidation with Active Directory

Microsoft believes that hub and spoke architectures represent the best approach to building comprehensive enterprise identity management solutions. To underscore this belief, Microsoft acquired Zoomit Corporation, the industry's leader in meta-directory technologies. Zoomit's flagship VIA product, renamed Microsoft Metadirectory Services

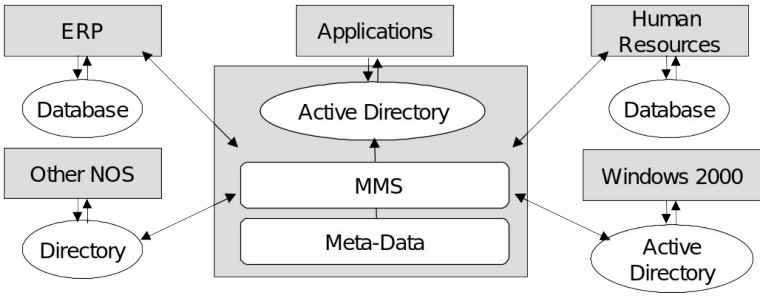

(MMS), implements a hub and spoke architecture.

MMS enables companies to use Active Directory to manage identity information stored in heterogeneous directory services. Microsoft also is enhancing MMS to become the industry's first comprehensive enterprise identity management platform by adding support for many popular applications and network services that store identity information in places other than directories.

Companies will see many benefits from MMS:

- It is easy to create a meta-directory, based on Active Directory, that applications, users, and administrators can use as a focal point for queries and updates.
- The brokering technologies in MMS enable companies to perform data transformation, track object locations, and maintain attribute ownership relationships across many different directories and repositories.
- Using MMS, companies can implement referential integrity protections and add value to the data stored in Active Directory.

Most important, Microsoft will continue to enhance MMS so that

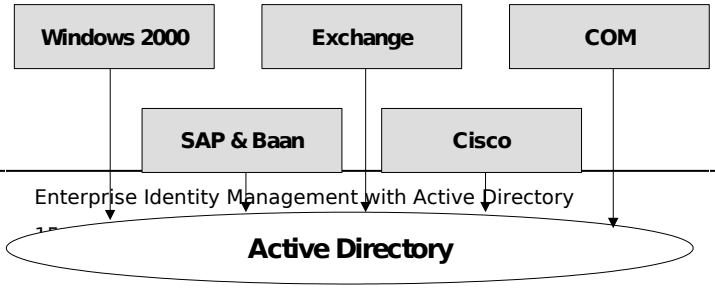

companies of all sizes will be able to realize the benefits of meta-directories and hub and spoke architectures.

Directory Consolidation

Microsoft's efforts around directory consolidation will focus on adapting Microsoft's products, and helping leading ISVs to adapt their products, to store identity data in Active Directory. The best example of an Active Directory-enabled Microsoft product is Exchange 2000, the next version of Microsoft Exchange Server. Exhange 2000 is the first version of Exchange Server that does not have its own directory service and will rely completely on Active Directory. Using Active Directory within Exchange 2000 enables customers to perform mailbox and user account administration in one place and eliminate the need to synchronize information with Windows 2000. Microsoft is working across other important product groups, such as the Microsoft SQL Server; Microsoft Office; Microsoft Site Server, Commerce Edition; and COM teams, to bring them similar benefits.

Microsoft is also working with leading ISVs to add Active Directory support to their products. For example, SAP and Baan have both demonstrated clients that locate services by searching for application identity information (in other words, the locations of servers) stored in Active Directory. Cisco Systems is extending a number of their products to use user identity information stored in Active Directory to deliver personalized networking services. In fact, because Microsoft believes so strongly that integration offers important benefits, Microsoft has added Active Directory integration requirements to the Windows 2000 Server logo specification. Customers who see the logo will know that an application is helping them achieve their goal of directory consolidation.

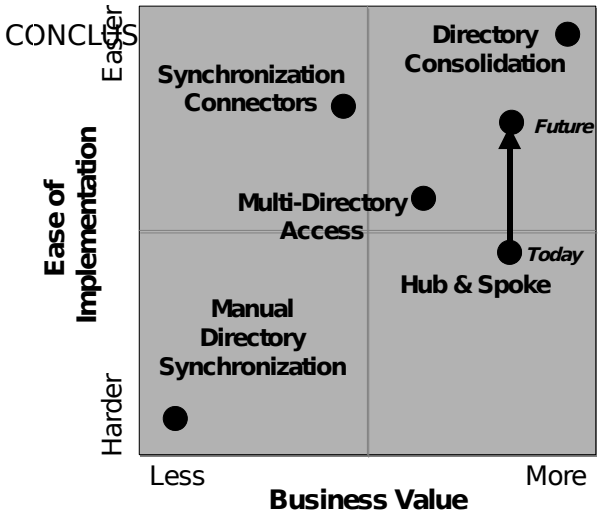

Given the complexity of enterprise identity management, most companies should plan on deploying a number of the different techniques outlined in this paper. However, it is important to note that each technique has a different cost versus business benefit ratio. Companies need to understand the tradeoffs shown on the grid above in order to maximize the benefits they gain over time.

Updating all directories and repositories manually clearly is the hardest approach to identity management and offers the least business value. Synchronization connectors offer more business value—especially in cases where synchronization requirements are simple—and are considerably less costly over time. Applications and scripts written using multi-directory access technologies can offer even more business value because they can contain customized business rules. Applications and scripts are harder to implement than connectors, however, because they require programming versus configuration.

Hub and spoke architectures offer considerable power and flexibility, and therefore the chance to solve significant identity management challenges. Today, however, they are harder to implement than multi-directory access solutions because developers need to understand specialized techniques used for maintaining data integrity and implementing business rules. Over time, hub and spoke architectures will become easier to implement as companies such as Microsoft roll out advanced capabilities such as templates, simplified rules languages, and pre-built connectors to popular applications.

Ultimately, the greatest return on investment comes from consolidating identity data around strategic directory services. This is because the number of directories where administrators must store and maintain identity data actually decreases. Synchronization is unnecessary. Here, the benchmark is how many applications support a given directory service instead of their own repositories.

Because of the depth and breadth of Microsoft’s solutions, Microsoft

believes that the Active Directory service of Windows 2000 is an ideal directory choice for whatever identity management path a company decides to take.